arXiv:1510.05354v1 [math.LO] 19 Oct 2015

RELATIVISED HOMOMORPHISM PRESERVATION AT THE FINITE LEVEL

LUCY HAM

Abstract. In this article, we investigate the status of the homomorphism preservation property amongst restricted classes of finite relational structures and algebraic structures. We show that there are many homomorphism-closed classes of finite lattices that are definable by a first-order sentence but not by existential positive sentences, demonstrating the failure of the homomorphism preservation property for lattices at the finite level. In contrast to the negative results for algebras, we establish a finite-level relativised homomorphism preservation theorem in the relational case. More specifically, we give a complete finite-level characterisation of first-order definable finitely generated anti-varieties relative to classes of relational structures definable by sentences of some general forms. When relativisation is dropped, this gives a fresh proof of Atserias's characterisation of first-order definable constraint satisfaction problems over a fixed template, a well known special case of Rossman's Finite Homomorphism Preservation Theorem.

1. Introduction

Finite-level preservation theorems: algebras vs relations. It is widely known that many classical results in model theory fail or become meaningless when considered at the finite level. An important class of such results are preservation theorems, which relate the form of syntactic expressions to the closure properties of classes they define. Classical preservation theorems are almost uniformly false when restricted to the class of all finite relational structures, but also tend to fail when restricted to the class all of finite algebras. A key example is the Łos-Tarski Preservation Theorem, which says that a first-order sentence is preserved under taking extensions if and only if it is equivalent to an existential sentence. See Tait [13] or Gurevich and Shelah [8] for a counterexample in the finite setting for relational structures, and Clark, Davey, Jackson and Pitkethly [5, Example 4.3] for a counterexample in the finite setting for algebras. For some preservation theorems only one side of the story is explained; for example the $\mathbb{ISP}$-Preservation Theorem, which asserts that a first-order formula is preserved under taking substructures and direct products if and only if it is equivalent to a universal-Horn sentence, remains to be settled at the finite-level in the relational setting. In contrast, Example 4.3 of [5] also gives a counterexample for algebraic structures at the finite level. On the other hand, Lyndon's Positivity Theorem, which says that a first-order sentence is preserved under taking surjective homomorphisms if and only if it is equivalent

Key words and phrases. Lattice, relational structure, constraint satisfaction problem, homomorphism preservation, antivariety.

The author would like to thank her PhD supervisors Brian Davey and Tomasz Kowalski, as well as Marcel Jackson for many useful discussions.

to a positive sentence, fails in the finite setting for relational structures (see Ajtai and Gurevich [1]), while its status remains unknown in the algebraic setting. A notable exception is Rossman's celebrated Finite Homomorphism Preservation Theorem [12, Theorem 5.16]: a first-order sentence is preserved under homomorphisms on finite relational structures if and only if it is equivalent, amongst finite relational structures, to an existential positive sentence. The question whether the classical Homomorphism Preservation Theorem would continue to hold on finite relational structures was a long-standing one in finite model theory. The analogous question for algebraic signatures remains unresolved. The proof for the relational case is highly non-trivial and, as Rossman suggests, the task of adjusting for function symbols is conceivably more difficult.

Relativised preservation theorems. Preservation theorems that hold for the class of all finite relational structures do not necessarily relativise to subclasses of this class, and moreover, while many classical preservation theorems fail amongst all finite relational structures, some preservation theorems are recovered in restricted classes of finite relational structures. In [3], Atserias, Dawar, and Grohe show that the Łos-Tarski Preservation Theorem is recovered in a number of classes of finite relational structures that are of interest from a computational point of view, including classes of finite acyclic graphs, classes of structures of bounded degree, and classes of structures of treewidth at most k, for any positive integer k. In [4], Atserias, Dawar and Kolaitis show that the homomorphism preservation property holds for a number of classes of finite relational structures, including classes of structures of bounded degree, classes of bounded treewidth, and more generally, classes of finite structures whose cores exclude at least one minor. However these results are not implied by Rossman's Finite Homomorphism Preservation Theorem. Furthermore, preservation theorems need not behave consistently when we restrict our attention to a particular subclass of the class of all finite relational structures: for example, the homomorphism preservation property holds for the class of all planar graphs, but in contrast, the extension preservation property fails in this setting.

Main results. Relativising preservation theorems to restricted classes of finite structures appears to be completely independent of whether or not the corresponding preservation property holds amongst all finite relational structures. On the other hand, relativisation of preservation theorems to restricted classes of finite algebras is not as well studied. In this paper, we are able to give some insight into this topic.

(1) We show that within the class of finite bounded lattices, there are many classes that demonstrate the failure of the homomorphism preservation property. More specifically, we show that for a fixed finite bounded lattice $\mathbf{L}$, the class of finite bounded lattices admitting a homomorphism into $\mathbf{L}$ is definable by a single first-order sentence, relative to the class of finite bounded lattices, but is not definable by universal sentences. See Theorem 3.1 in Section 3.
(2) In contrast to (1), we establish a relativised homomorphism preservation theorem (see Theorem 4.8 in Section 4) for a wide range of classes of finite relational structures. Theorem 4.8 gives us a complete characterisation of

first-order definable classes of relational structures admitting a homomorphism into a fixed finite relational structure, relative to classes $\mathcal{K}$ definable by sentences of some general forms. In Theorem 4.16 we extend this characterisation to classes of relational structures admitting a homomorphism into one of a finite number of relational structures.

We will see that Theorem 3.1 mentioned in (1) not only demonstrates the failure of the homomorphism preservation property for the class of finite bounded lattices, but also the failure of the extension preservation property (Łos-Tarski Preservation Theorem). So the extension preservation property, which fails for the class of all finite algebras and also the class of all finite relational structures, continues to fail when we restrict to the class of finite bounded lattices. This result could be contrasted with the positive results of Atserias, Dawar and Grohe [3] mentioned previously, which show that the extension preservation property is recovered in a number of restricted classes of finite relational structures.

When the class $\mathcal{K}$ in Theorem 4.8 is relaxed to include all finite relational structures, our result coincides with a well known one due to Atserias, which says that first-order definable constraint satisfaction problems are precisely those with finite tree duality (or equivalently, those definable by existential positive sentences), see Theorem 9 of [2]. It is also a consequence Rossman's Finite Homomorphism Preservation Theorem, which holds for classes of relational structures admitting a homomorphism into one of *any* number of finite relational structures. Our proof, however, is independent of both Atserias and Rossman and relies mostly on some constructions borrowed from Larose, Loten and Tardif given in [10]. We remark that while it remains unclear if the homomorphism preservation property in its full generality holds in these classes, the restricted preservation property, nevertheless, covers many commonly encountered homomorphism-closed classes.

Article structure. This article is comprised of three main sections. We begin with some background and definitions that will be required for the remainder of the paper, we then investigate the status of Relativised Homomorphism Preservation Theorem 2.13 at the finite level for both algebraic and relational signatures. In Section 3, we give counterexamples to the finite-level version of Relativised Homomorphism Preservation Theorem 2.13 for algebraic structures. Moreover, we completely characterise those classes of finite lattices admitting a homomorphism into a fixed finite bounded lattice. In Section 4, we establish a restricted form of the finite-level version of Relativised Homomorphism Theorem 2.13 for relational structures. The proof follows an Ehrenfeucht-Fraïssé game argument on structures obtained by adjusting some constructions borrowed from Larose, Loten and Tardif given in [10].

2. Preliminaries

Definition 2.1. A *signature* (or *language* or *type*) is a pair $\mathcal{L} = \langle \mathcal{R}, \mathcal{F} \rangle$ consisting of a set of relation symbols $\mathcal{R}$ and a set of function symbols $\mathcal{F}$. Each relation and function symbol has an associated *arity* n; we do not allow for nullary relation symbols so n is quantified over $\mathbb{N}$ in this case; we do however include the possibility of 0-ary function symbols so n is quantified over $\mathbb{N} \cup \{0\}$ in this case. An $\mathcal{L}$-*structure* $\mathbf{A} = \langle A; \mathcal{R}^{\mathbf{A}}, \mathcal{F}^{\mathbf{A}} \rangle$ consists of a universe A together with an interpretation of

- each n-ary relation symbol $R \in \mathcal{R}$ as a subset $R^{\mathbf{A}}$ of A^n,

- each n-ary function symbol $F \in \mathcal{F}$ as a function $F^{\mathbf{A}} \colon A^n \to A$.

If $\mathcal{R} = \varnothing$, then $\mathbf{A}$ is an *algebra*, if $\mathcal{F} = \varnothing$ then $\mathbf{A}$ is a *relational structure*.

Definition 2.2. Let $\mathbf{A}$ be a structure in some signature $\mathcal{L}$ consisting of possibly both relation and function symbols. We let $\mathcal{L}_{\mathbf{A}}$ be the language obtained by adding a nullary function symbol a to $\mathcal{L}$, for each $a \in A$.

Definition 2.3. Let $n \in \mathbb{N}$, let $\mathcal{R} = \{R_1, \dots, R_n\}$ be a purely relational signature, and let $\mathbf{A} = \langle A; R_1^{\mathbf{A}}, \dots, R_n^{\mathbf{A}} \rangle$. A *weak substructure* of $\mathbf{A}$ is a structure $\mathbf{A'} = \langle A'; R_1^{\mathbf{A}'}, \dots, R_n^{\mathbf{A}'} \rangle$, where $A' \subseteq A$ and $R_i^{\mathbf{A}'} \subseteq R_i^{\mathbf{A}}$ for each $i \in \{1, \dots, n\}$. The weak substructure $\mathbf{A'}$ is *proper* if either $A' \subsetneqq A$ or $A' = A$ and there exists $i \leq n$ such that $R_i^{\mathbf{A}'} \subsetneqq R_i^{\mathbf{A}}$.

Definition 2.4. Let $\mathcal{L}$ be some signature consisting of possibly both relation and function symbols, and let $\mathcal{J}$ be a class of $\mathcal{L}$-structures. We define $\mathbb{H}^{\leftarrow}(\mathcal{J})$ to be the class of all $\mathcal{L}$-structures admitting a homomorphism into some member of the class $\mathcal{J}$.

Definition 2.5. Let I be a set and let $\mathcal{F} = \{\mathbf{A}_i \mid i \in I\}$ be a set of structures in some signature $\mathcal{L}$. We say that $\mathcal{F}$ is a *homomorphism-independent* set if for all $i, j \in I$ with $i \neq j$, there is no homomorphism from $\mathbf{A}_i$ to $\mathbf{A}_j$.

Definition 2.6. Let $n, m \in \mathbb{N}$ and let $\mathcal{F} = \{\mathbf{A}_1, \dots \mathbf{A}_n\}$ be a finite set of finite structures in the purely relational signature $\mathcal{R} = \{R_1, \dots, R_m\}$. We denote the finite models in $\mathbb{H}^{\leftarrow}(\mathcal{F})$ by $\mathbb{H}^{\leftarrow}_{\text{fin}}(\mathcal{F})$. An *obstruction* for the class $\mathcal{F}$ is any finite $\mathcal{R}$-structure not belonging to $\mathbb{H}^{\leftarrow}_{\text{fin}}(\mathcal{F})$. An obstruction $\mathbf{C}$ for the class $\mathcal{F}$ is *critical* if every proper weak substructure of $\mathbf{C}$ belongs to $\mathbb{H}^{\leftarrow}_{\text{fin}}(\mathcal{F})$. We say that $\mathbb{H}^{\leftarrow}_{\text{fin}}(\mathcal{F})$ satisfies the *finite duality property* if there exists a finite set $\mathcal{S}$ of finite structures with the following property: a finite $\mathcal{R}$-structure $\mathbf{B}$ belongs to the class $\mathbb{H}^{\leftarrow}_{\text{fin}}(\mathcal{F})$ if and only if $\mathbf{B}$ avoids a homomorphism from every member of $\mathcal{S}$. Let $\mathbf{A}$ be a finite structure in the purely relational signature $\mathcal{R}$.

The class of all finite $\mathcal{R}$-structures admitting a homomorphism into $\mathbf{A}$ is denoted $\text{CSP}(\mathbf{A})$. The membership problem for $\text{CSP}(\mathbf{A})$ is known as the *constraint satisfaction problem* over template $\mathbf{A}$. Note that $\mathbb{H}^{\leftarrow}_{\text{fin}}(\mathcal{F}) = \bigcup_{i \in \{1,\dots,n\}} \text{CSP}(\mathbf{A}_i)$.

Definition 2.7. Let $m \in \mathbb{N}$, let $\mathbf{A}$ be a finite relational structure in the purely relational signature $\mathcal{R} = \{R_1, \dots, R_m\}$ with R_i of arity r_i, for all $i \in \{1, \dots, m\}$. The *incidence multigraph* of $\mathbf{A}$, denoted $\text{Inc}(\mathbf{A})$, is defined as the bipartite multigraph with parts A and $\text{Block}(\mathbf{A})$, where $\text{Block}(\mathbf{A})$ consists of all pairs $(R_i, (a_1, \dots, a_{r_i}))$ such that $R_i \in \{R_1, \dots R_m\}$ and $(a_1, \dots, a_{r_i}) \in R_i^{\mathbf{A}}$, and with edges $e_{a,k,B}$ joining $a \in A$ to $B = (R_j, (a_1, \dots, a_{r_j})) \in \text{Block}(\mathbf{A})$, when $a_k = a$ for some $k \in \{1, \dots, r_j\}$. Let a and b be two elements in A. Then the *distance between a and b in* $\mathbf{A}$ $\text{dist}_{\mathbf{A}}(a, b)$ is defined to be half the number of edges in a shortest path connecting them in $\text{Inc}(\mathbf{A})$. The *girth* of $\mathbf{A}$ is defined to be half the shortest length of a cycle in $\text{Inc}(\mathbf{A})$. The *diameter* of $\mathbf{A}$ is defined to be $\max\{\text{dist}_{\mathbf{A}}(u, v) \mid (u, v) \in A^2\}$.

Definition 2.8. Let $\mathcal{K}$ be a class of $\mathcal{L}$-structures and let $\mathcal{J}$ be a subclass of $\mathcal{K}$. We define the following class:

$$\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathcal{J}) := \{\mathbf{B} \in \mathcal{K} \mid \text{there exists } \mathbf{J} \in \mathcal{J} \text{ and a homomorphism from } \mathbf{B} \text{ to } \mathbf{J}\}.$$

In the case where $\mathcal{K}$ is the class of all $\mathcal{L}$-structures we abbreviate $\mathbb{H}^{\leftarrow}_{\mathcal{K}}$ to $\mathbb{H}^{\leftarrow}$. In general, $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathcal{J}) = \mathbb{H}^{\leftarrow}(\mathcal{J}) \cap \mathcal{K}$.

Let $n \in \mathbb{N}$, let $\mathbf{A} \in \mathcal{K}$, and let $\mathcal{F} = \{\mathbf{A}_1, \ldots, \mathbf{A}_n\}$ be a finite set of structures in $\mathcal{K}$. We follow Gorbunov [7] and refer to the class $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathbf{A})$ as a *principal colour-family in* $\mathcal{K}$, and we refer to the class $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathcal{F})$ as a *colour-family in* $\mathcal{K}$.

Definition 2.9. Let $n \in \mathbb{N} \cup \{0\}$ and let $m \in \mathbb{N}$. A sentence ψ is called an *anti-identity* if it is of the form

$$\forall x_1, \ldots, x_n \, [\neg\alpha_1(x_1, \ldots, x_n) \vee \cdots \vee \neg\alpha_m(x_1, \ldots, x_n)],$$

where $\alpha_i(x_1, \ldots, x_n)$ is an atomic formula for all $i \in \{1, \ldots, m\}$.

Notation 2.10. Let Σ be a set of first-order sentences in some signature $\mathcal{L}$. We let $\mathrm{Mod}(\Sigma)$ denote the class of all models of Σ.

.

Definition 2.11. A subclass $\mathcal{K}'$ of a class $\mathcal{K}$ is called *strictly elementary relative to* $\mathcal{K}$ if there is a first-order sentence ϕ such that $\mathcal{K}' = \mathcal{K} \cap \mathrm{Mod}(\phi)$.

Definition 2.12. A subclass $\mathcal{K}'$ of a class $\mathcal{K}$ is called an *anti-variety relative to* $\mathcal{K}$ if $\mathcal{K}' = \mathcal{K} \cap \mathrm{Mod}(\Sigma)$ where Σ is a set of anti-identities.

The proof of the Relativised Homomorphism Preservation Theorem 2.13 requires only standard applications of some model-theoretic fundamentals. We have included its proof, however, for completeness sake. We remark that the results of this paper are more easily formulated for classes that are complements of homomorphism closed classes. Such classes are closed under taking homomorphic pre-images: classes $\mathcal{J}$ satisfying $\mathbf{B} \in \mathcal{J}$ whenever $\mathbf{B}$ admits a homomorphism to $\mathbf{A}$ and $\mathbf{A} \in \mathcal{J}$. For this reason, we work with the homomorphism preservation theorem in its complemented form.

Theorem 2.13 (Relativised Homomorphism Preservation Theorem)**.** *Let Σ be a set of sentences in the signature $\mathcal{L}$, let $\mathcal{K} = \mathrm{Mod}(\Sigma)$, and let $\mathcal{J}$ be a subclass of $\mathcal{K}$. Then the class $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathcal{J})$ is an anti-variety relative to $\mathcal{K}$. Moreover, if $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathcal{J})$ is strictly elementary with respect to $\mathcal{K}$, then there is a sentence consisting of a finite conjunction anti-identities that is logically equivalent to ϕ modulo Σ.*

Proof. We first show that $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathcal{J})$ is axiomatisable by a set of anti-identities with respect to the class $\mathcal{K}$. The second part of the theorem will follow by compactness.

Let $\mathcal{J}^* = \mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathcal{J})$ and let $\mathrm{Th}_{\mathrm{anti}}(\mathcal{J})$ be the set of anti-identities true in $\mathcal{J}$. Clearly, if $\mathbf{A} \in \mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathcal{J})$, then $\mathbf{A} \models \mathrm{Th}_{\mathrm{anti}}(\mathcal{J})$, since anti-identities are preserved under taking homomorphic pre-images. Now let $\mathbf{A} \in \mathcal{K}$ and suppose that $\mathbf{A} \models \mathrm{Th}_{\mathrm{anti}}(\mathcal{J})$. Let $\mathbf{T}_{\mathbf{A}}$ be the set of atomic sentences true in $\mathbf{A}$ in the expanded signature $\mathcal{L}_{\mathbf{A}}$, and let $\mathcal{J}^*_{\mathbf{A}}$ be the class of all $\mathcal{L}_{\mathbf{A}}$-structures whose $\mathcal{L}$-reduct belongs to $\mathcal{J}^*$. For any finite subset $\{\phi_1(a_1, \ldots, a_n), \ldots, \phi_k(a_1, \ldots, a_n)\}$ of $\mathbf{T}_{\mathbf{A}}$, we have

$$\mathbf{A} \models \phi_1(a_1, \ldots, a_n) \wedge \cdots \wedge \phi_k(a_1, \ldots, a_n).$$

So, $\mathbf{A}$ fails the sentence

$$\sigma \equiv \forall x_1, \ldots, x_n \, [\neg\phi_1(x_1, \ldots, x_n) \vee \cdots \vee \neg\phi_k(x_1, \ldots, x_n)].$$

Hence there is some structure $\mathbf{B} \in \mathcal{J}^*$ that fails σ.

Let $b_1, \ldots, b_n$ be the elements of B that witness the failure of the sentence σ. Now expand $\mathbf{B}$ in the signature $\mathcal{L}_{\mathbf{A}}$ under the following interpretation: let b be any

element of B, and define $a_i^{\mathbf{B}} := b_i$, if $i \in \{1, \dots, n\}$, and $a_i^{\mathbf{B}} := b$, otherwise. Then $\mathbf{B}_{\mathbf{A}} \in \mathcal{J}^*_{\mathbf{A}}$, and satisfies

$$\phi_1(a_1, \dots, a_n) \wedge \cdots \wedge \phi_k(a_1, \dots, a_n).$$

Since $\mathcal{K}$ is elementary and $\mathcal{J}^*$ is first order definable relative to $\mathcal{K}$ it follows, by compactness, that there is some $\mathbf{Y}_{\mathbf{A}} \in \mathcal{J}^*_{\mathbf{A}}$ such that $\mathbf{Y}_{\mathbf{A}} \models \mathbf{T}_{\mathbf{A}}$. Now let $f : A \to Y$ be the map that sends each element a_i of $\mathbf{A}$ to its corresponding interpretation in $\mathbf{Y}_{\mathbf{A}}$. Then, for every atomic formula $\phi(x_1, \dots, x_n)$ in $\mathcal{L}$ and every n-tuple $(a_1, \dots, a_n) \in A^n$, we have

$$\mathbf{A} \models \phi(a_1, \dots, a_n) \Longrightarrow \mathbf{Y} \models \phi(f(a_1), \dots f(a_n)).$$

Thus, the map $f : A \to Y$ is a homomorphism from $\mathbf{A}$ to $\mathbf{Y}$. Since $\mathbf{Y} \in \mathcal{J}^*$, there exists a homomorphism g from $\mathbf{Y}$ to some member $\mathbf{J}$ of $\mathcal{J}$, and so $g \circ f$ is a homomorphism from $\mathbf{A}$ to $\mathbf{J}$, thus $\mathbf{A} \in \mathcal{J}^*$. Hence we have shown that there is a set Φ of anti-identities such that $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathcal{J}) = \mathcal{K} \cap \mathrm{Mod}(\Phi)$, namely $\Phi = \mathrm{Th}_{\mathrm{anti}}(\mathcal{J})$. Now assume there is a single first-order sentence ϕ such that $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathcal{J}) = \mathcal{K} \cap \mathrm{Mod}(\phi)$. Then every model of $\Sigma \cup \Phi$ is also a model of $\Sigma \cup \{\phi\}$. By the Completeness Theorem of first-order logic, we get that $\Sigma \cup \Phi \vdash \{\phi\}$. It then follows, by compactness, that there is a finite (possibly empty) set of anti-identities $F \subseteq \Phi$ such $\Sigma \cup F \vdash \phi$. Thus, the sentence ϕ is logically equivalent to the conjunction of anti-identities $\bigwedge F$ modulo Σ. □

3. Failure at the finite level

We now show that the finite-level version of the Relativised Homomorphism Preservation Theorem 2.13 fails for algebraic signatures. If we restrict the class $\mathcal{K}$ to all finite bounded lattices, then we can find many classes closed under taking homomorphic pre-images that are first-order definable relative to $\mathcal{K}$, but not definable by universal sentences.

Theorem 3.1. *Let $\mathcal{K}$ be the class of finite bounded lattices in the signature $\mathcal{L} = \{\vee, \wedge, 0, 1\}$ and let $\mathbf{L}$ be any lattice in $\mathcal{K}$ with $|L| \geqslant 2$. The class $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathbf{L})$ of all finite bounded lattices admitting a homomorphism into $\mathbf{L}$ is strictly first-order definable relative to $\mathcal{K}$ but not definable by any universal sentence.*

Proof. The class $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathbf{L})$ of all finite bounded lattices admitting a homomorphism into $\mathbf{L}$ is definable by a $\forall\exists$ sentence ϕ by [5, Theorem 4.2]. We will show that the class $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathbf{L})$ is not definable by any universal sentence. Specifically, we show that, for each $n \in \mathbb{N}$, there exists a finite bounded lattice $\mathbf{A}_n$ satisfying the following two conditions.

(1) $\mathbf{A}_n \notin \mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathbf{L})$;
(2) every n-generated sublattice of $\mathbf{A}_n$ belongs to $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathbf{L})$.

For each $n \in \mathbb{N}$, let $\mathbf{A}_n$ be the lattice depicted in Figure 1. The lattice $\mathbf{A}_n$ is a "stack" of n copies of $\mathbf{M}_k$, where k is a natural number strictly greater than $\max\{|L|, 2\}$.

(1): Let $n \in \mathbb{N}$. The smallest lattice that $\mathbf{A}_n$ collapses to, without identifying 0 and 1, is isomorphic to the lattice $\mathbf{M}_k$. Since $\mathbf{M}_k$ is simple (as $k \geqslant 3$), every congruence on $\mathbf{M}_k$ collapses everything or collapses nothing. Thus if $\mathbf{A}_n$ is to map homomorphically into $\mathbf{L}$, then $\mathbf{L}$ must contain a sublattice isomorphic to $\mathbf{M}_k$, but this is not possible since $\mathbf{M}_k$ contains more elements than $\mathbf{L}$. Hence $\mathbf{A}_n \notin \mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathbf{L})$.

(2): We will show that every n-generated sublattice of $\mathbf{A}_n$ has a homomorphism onto the two-element lattice $\mathbb{2}$, and therefore has a homomorphism into $\mathbf{L}$. Assume that we are selecting n generators from the lattice $\mathbf{A}_n$. In order to generate a sublattice of $\mathbf{A}_n$ that does not admit a homomorphism onto $\mathbb{2}$ (or equivalently, a sublattice that does not contain a prime ideal), we would need to select *at least* three elements from each of the n maximal anti-chains in $\mathbf{A}_n$. Thus, the lattice $\mathbf{P}$ given in Figure 2 is the smallest sublattice of $\mathbf{A}_n$ that does not admit a homomorphism onto $\mathbb{2}$. But this lattice is $3n$-generated and not generated by fewer than $3n$ elements. Hence every n-generated a sublattice of $\mathbf{A}_n$ admits a homomorphism onto $\mathbb{2}$, and therefore into $\mathbf{L}$. □

Remark 3.2. Notice that within the class of finite bounded lattices, the (first-order definable) property of admitting a homomorphism onto $\mathbb{2}$ is preserved by taking sublattices, and so the above examples also give us a counterexample to the finite-level relativised Łos-Tarski Preservation Theorem for algebraic signatures.

Remark 3.3. Also notice that the proof of Theorem 3.1 works for the class $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathcal{L})$, where $\mathcal{L}$ is any class of finite bounded lattices containing at least one non-trivial lattice, and for which there exists $k \geqslant 3$ such that $\mathbf{M}_k \notin \mathbb{S}(\mathcal{L})$.

Remark 3.4. The proof of Theorem 3.1 also works with $\mathcal{K}$ replaced by any class of finite lattices containing the lattices in Figure 1 (for each $n \in \mathbb{N}$). As the lattice $\mathbf{A}_n$ is modular, the class $\mathcal{K}$ in Theorem 3.1 could be taken to be the class of all finite bounded modular lattices. Alternatively, $\mathcal{K}$ could to be the finite part of the variety generated by the bounded lattices $\mathbf{A}_n$. This variety can be shown to coincide with the variety generated by $\mathbf{M}_\omega$.

4. A Relativised Homomorphism Preservation Theorem

We now turn to the purely relational side of things. In this section, we establish a finite-level version of the Relativised Homomorphism Preservation Theorem for relational signatures. We present this result as Theorem 4.8 below. We first prove the non-relativised version of Theorem 4.8, namely Theorem 4.9, from which Theorem 4.8 will follow almost immediately. The proof requires three constructions given in Definition 4.1, Definition 4.2 and Definition 4.3 below, borrowed from [10]. The quotient structure given in Definition 4.3, will play a key role in establishing Theorem 4.8. We follow Jackson and Trotta [9] and refer to this quotient structure as the *n-pinch* over $\mathbf{A}$. In [10], Larose, Loten and Tardif show that, for a relational structure $\mathbf{A}$, demonstrating membership of the n-pinch in $\mathrm{CSP}(\mathbf{A})$, for some natural number n, is equivalent to $\mathrm{CSP}(\mathbf{A})$ admitting the finite duality property. This fact will be crucial in the proof of Theorem 4.8.

We conclude the section by extending Theorem 4.8. More specifically, we obtain a complete characterisation of not only first-order definable principal colour-families but also first-order definable colour-families in classes $\mathcal{K}$ closed under forming the n-pinch construction and taking disjoint unions; see Theorem 4.16.

In this section, we assume that all structures are finite in a finite purely relational signature $\mathcal{R} = \{R_1, \ldots, R_m\}$ with $m \in \mathbf{N} \cup \{0\}$ and R_i of arity r_i, for all $i \in \{1, \ldots, m\}$, unless otherwise stated.

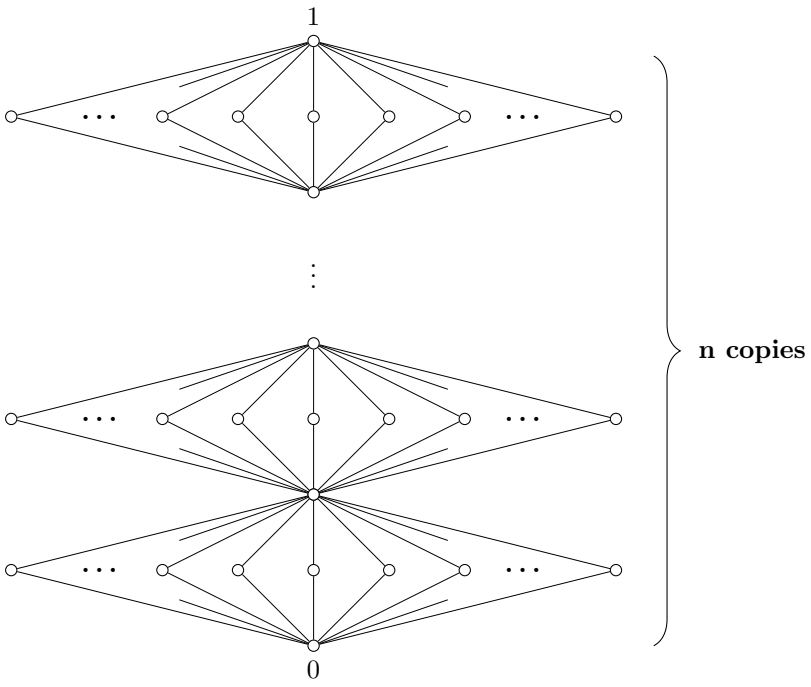

FIGURE 1. The lattice $\mathbf{A}_n$.

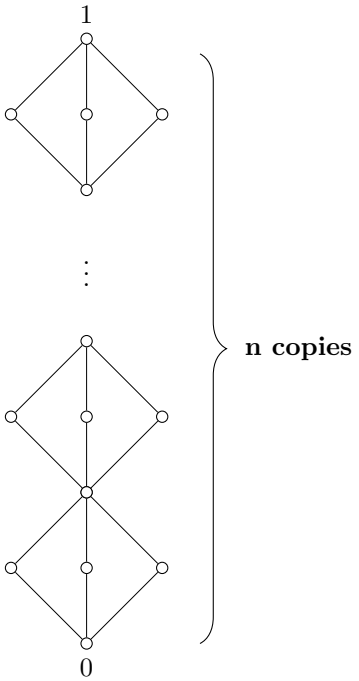

FIGURE 2. The sublattice $\mathbf{P}$ of $\mathbf{A}_n$.

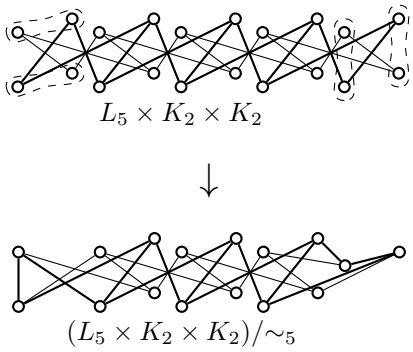

FIGURE 3. The structure $\mathbf{P}_5(\mathbf{K}_2)$ as a quotient of $L_5 \times K_2 \times K_2$.

Definition 4.1. Let $n, m \in \mathbb{N} \cup \{0\}$, the *n-link* of signature $\mathcal{R} = \{R_1, \ldots, R_m\}$ is defined to be the structure

$$\mathbf{L}_n = \langle \{0, 1, \ldots, n\}; {R_1}^{\mathbf{L}_n}, \ldots, {R_m}^{\mathbf{L}_n} \rangle,$$

where ${R_i}^{\mathbf{L}_n} = \bigcup_{j=1}^{n} \{j-1, j\}^{r_i}$, for all $i \in \{1, \ldots, m\}$.

Definition 4.2. Let $n \in \mathbb{N}$ and let $\mathbf{A}$ be a structure in the signature $\mathcal{R} = \{R_1, \ldots, R_m\}$. Consider the product structure $\mathbf{L}_n \times \mathbf{A} \times \mathbf{A}$. We define an equivalence relation $\sim_n$ on $\mathbf{L}_n \times \mathbf{A} \times \mathbf{A}$ as follows:

$$(i, a, b) \sim_n (i', a', b') \iff \begin{cases} (i, a, b) = (i', a', b'), & \text{or} \\ i = i' = 0 \quad \text{and} \quad a = a', & \text{or} \\ i = i' = n \quad \text{and} \quad b = b'. \end{cases}$$

Now using Definition 4.1 and Definition 4.2 we construct the following quotient structure.

Definition 4.3. Let $n \in \mathbb{N}$ and let $\mathbf{A}$ be a structure in the signature $\mathcal{R} = \{R_1, \ldots, R_m\}$, the n**-pinch** over $\mathbf{A}$ is defined to be the structure

$$\mathbf{P}_n(\mathbf{A}) = (\mathbf{L}_n \times \mathbf{A} \times \mathbf{A})/\sim_n.$$

Observe that the first coordinate of the elements in $\mathbf{L}_n \times \mathbf{A} \times \mathbf{A}$ is constant on each of the $\sim_n$-equivalence classes. In other words, the first coordinate is unchanged by the formation of the quotient structure $\mathbf{P}_n(\mathbf{A})$. We therefore use the first coordinate as a means for specifying the position of an element in a copy of $\mathbf{A}$ or $\mathbf{A} \times \mathbf{A}$ in the n-pinch: notationally, we define a function $\iota : \mathbf{P}_n(\mathbf{A}) \to \{0, \ldots, n\}$ by $\iota((i, a, b)/\sim_n) := i$. We refer to the copy of $\mathbf{A}$ in $\mathbf{P}_n(\mathbf{A})$ at position $i = 0$ as the *left pinch* of $\mathbf{P}_n(\mathbf{A})$, and we refer to the copy of $\mathbf{A}$ at position $i = n$ as the *right pinch* of $\mathbf{P}_n(\mathbf{A})$. We denote the substructures of $\mathbf{P}_n(\mathbf{A})$ induced by the sets $B_0 = \{(k, a, b)/\sim_n |\ k \neq 0\}$ and $B_n = \{(k, a, b)/\sim_n |\ k \neq n\}$ as $\mathbf{B}_R$ and $\mathbf{B}_L$, respectively. We refer to the copy of $\mathbf{A} \times \mathbf{A}$ in $\mathbf{B}_R$ at position $i = 1$ as the *right slice* of $\mathbf{B}_R$, and we refer to the copy of $\mathbf{A} \times \mathbf{A}$ in $\mathbf{B}_L$ at position $i = n - 1$ as the *left slice* of $\mathbf{B}_L$.

Example 4.4. *Let* $\mathbf{K}_2$ *be the complete graph on two vertices. The formation of* $\mathbf{P}_5(\mathbf{K}_2)$ *is depicted in Figure 3. The graph of* $\mathbf{P}_5(\mathbf{K}_2)$ *can be visualised as a Möbius strip.*

We now collect together some facts that will be used in the proof of our main result.

Theorem 4.5. *Let* $\mathbf{A} = \langle A; R_1^{\mathbf{A}}, \dots, R_m^{\mathbf{A}} \rangle$ *be a finite relational structure.*

(1) ([10]*, Lemma* 4.6) *For each natural number* n*, the substructures* $\mathbf{B}_R$ *and* $\mathbf{B}_L$ *of* $\mathbf{P}_n(\mathbf{A})$ *admit homomorphisms into* $\mathbf{A}$*.*
(2) ([10] *Theorem* 2.5 *and Theorem* 4.7) *The following are equivalent*:
   - CSP($\mathbf{A}$) *satisfies the finite duality property*;
   - *there exists some natural number* n *such that* $\mathbf{P}_n(\mathbf{A})$ *admits a homomorphism to* $\mathbf{A}$*.*
(3) ([9] *Lemma* 21) *For each* $n \in \mathbb{N}$*, there exists a surjective homomorphism from* $\mathbf{P}_{n+1}(\mathbf{A})$ *to* $\mathbf{P}_n(\mathbf{A})$*.*

Remark 4.6. For (1), it is routine to check that the map $\phi : B_R \to A$ defined by $\phi((k, a, b)/\!\sim_n) = b$ is a homomorphism from $\mathbf{B}_R$ to $\mathbf{A}$ and similarly, that the map $\varphi : B_L \to A$ defined by $\phi((k, a, b)/\!\sim_n) = a$ is a homomorphism from $\mathbf{B}_R$ to $\mathbf{A}$, see the proof of Lemma 4.6 in [10]. The proof of (2) depends mostly on the observation that there are only finitely many critical obstructions of a given diameter (see Lemma 2.4 of [10]. For (3), the map defined by

$$\phi_n((i, a, b)/\!\sim_n) = \begin{cases} (i, a, b)/\!\sim_n, & \text{if } i \leq n, \\ (n, a, b)/\!\sim_n, & \text{if } i = n + 1. \end{cases}$$

is easily shown to be a surjective homomorphism, see the proof of Lemma 21 in [10].

Definition 4.7. Let $\mathcal{K}$ be a class of $\mathcal{R}$-structures and let $\mathcal{F}$ be a subclass of $\mathcal{K}$. We say that $\mathcal{K}$ is *closed under forming the* n*-pinch construction over* $\mathcal{F}$, if for every $n \in \mathbb{N}$ and every $\mathbf{F} \in \mathcal{F}$, we have that $\mathbf{P}_n(\mathbf{F})$ and its substructures $\mathbf{B}_R$ and $\mathbf{B}_L$ belong to the class $\mathcal{K}$.

We now state the main result for this section: a finite analogue of the Relativised Homomorphism Theorem 2.13 in the case where $\mathcal{J}$ is a single structure.

Theorem 4.8. *Let* $\mathcal{K}$ *be a class of finite* $\mathcal{R}$*-structures and let* $\mathbf{A} \in \mathcal{K}$*. If* $\mathcal{K}$ *is closed under forming the* n*-pinch construction over* $\mathbf{A}$ *and taking disjoint unions, then* $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathbf{A})$ *is equal to* $\mathcal{K} \cap \mathrm{Mod}(\phi)$ *for some first order sentence* ϕ *if and only if there is a finite conjunction of anti-identities* ψ *such that* $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathbf{A}) = \mathcal{K} \cap \mathrm{Mod}(\psi)$*.*

We first prove a result from which Theorem 4.8 will follow almost immediately.

Theorem 4.9. *Let* $\mathbf{A}$ *be a finite relational structure in the signature* $\mathcal{R}$*. Then the following are equivalent:*

(1) CSP($\mathbf{A}$) *satisfies the finite duality property*;
(2) CSP($\mathbf{A}$) *is definable by a finite conjunction of anti-identities*;
(3) CSP($\mathbf{A}$) *is first-order definable*;

Proof. The implication (1) $\Rightarrow$ (2) is well known (see Lemma 2 in [9] for example), and (2) $\Rightarrow$ (3) is trivial.

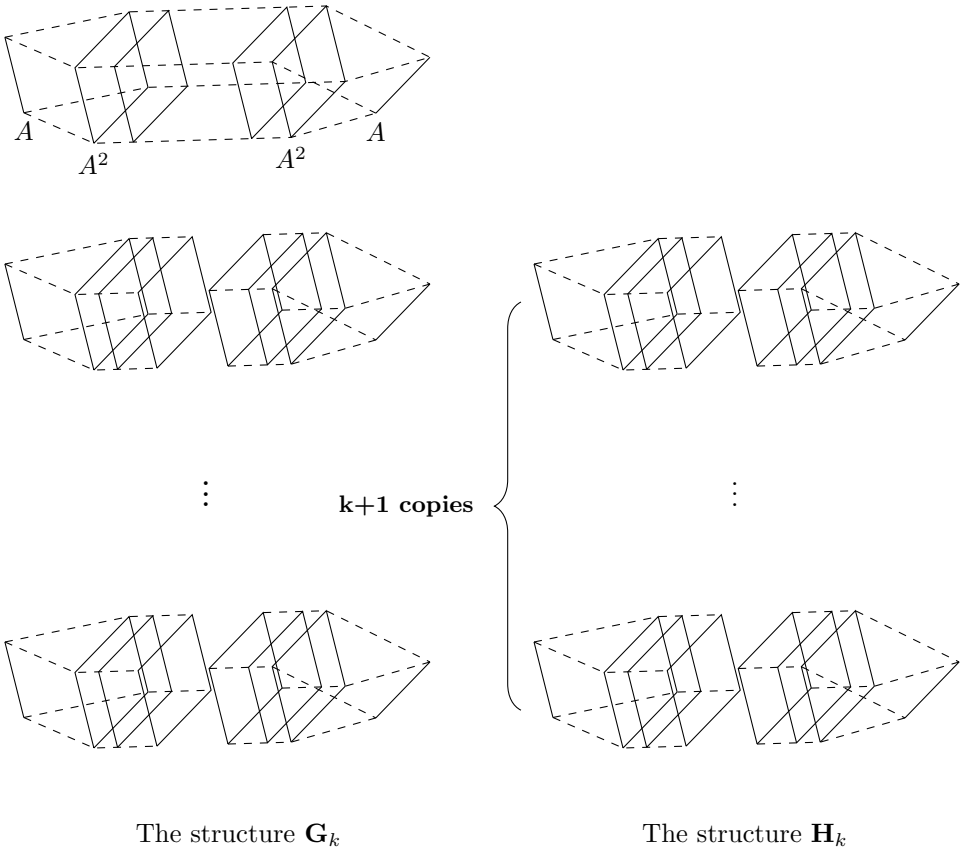

FIGURE 4. The Ehrenfeucht-Fraïssé game "board".

We now prove (3) ⇒ (1) by establishing the contrapositive. We show that for a finite relational structure $\mathbf{A}$, if the class CSP($\mathbf{A}$) fails the finite duality property, then the property of admitting a homomorphism to $\mathbf{A}$ is not expressible by a first-order sentence. The proof follows an Ehrenfeucht-Fraïssé game arguement; we show that, for each non-negative integer k, there are two structures $\mathbf{G}_k$ and $\mathbf{H}_k$ such that:

- Duplicator has a winning strategy in the k-round Ehrenfeucht-Fraïssé game played on $\mathbf{G}_k$ and $\mathbf{H}_k$, and
- $\mathbf{H}_k$ belongs to CSP($\mathbf{A}$), but $\mathbf{G}_k$ does not.

If CSP($\mathbf{A}$) does not satisfy the finite duality property, then Theorem 4.5 tells us that for each natural number n, there exists no homomorphism from $\mathbf{P}_n(\mathbf{A})$ to $\mathbf{A}$, but the substructures $\mathbf{B}_R$ and $\mathbf{B}_L$ of $\mathbf{P}_n(\mathbf{A})$ admit homomorphisms to $\mathbf{A}$. This suggests that we construct our structures $\mathbf{G}_k$ and $\mathbf{H}_k$ using combinations of $\mathbf{P}_n(\mathbf{A})$ and its substructures $\mathbf{B}_R$ and $\mathbf{B}_L$, where n is some number defined in terms of k.

Let k be a non-negative integer and assume that n' is an integer strictly greater than $2^{k+1}+1$. Now consider $\mathbf{P}_{n'}(\mathbf{A})$ and the disjoint union $\mathbf{B}_R \dot\cup \mathbf{B}_L$. We build $\mathbf{G}_k$ and $\mathbf{H}_k$ in the following way. Let $\mathbf{H}_k$ be the structure obtained by taking the disjoint union of $k+1$ copies of $\mathbf{B}_R \dot\cup \mathbf{B}_L$, and let $\mathbf{G}_k$ be $\mathbf{P}_{n'} \dot\cup \mathbf{H}_k$. See Figure 4 for an abstract depiction of the structures $\mathbf{G}_k$ and $\mathbf{H}_k$.

We begin by defining a notion of what it means for two elements in $\mathbf{G}_k$ or $\mathbf{H}_k$ to be "close to" or "far away" from each other; but first we need to give a concrete definition of distance between two elements in $\mathbf{G}_k$ or $\mathbf{H}_k$. We will refer to a copy of $\mathbf{P}_{n'}(\mathbf{A})$ or $\mathbf{B}_{\mathrm{R}}$ or $\mathbf{B}_{\mathrm{L}}$ in $\mathbf{G}_k$ or $\mathbf{H}_k$ as a *block-component*. Every element x in $\mathbf{G}_k$ or $\mathbf{H}_k$ is a copy of one of the form $(i, a, b)/{\sim_{n'}}$, thus we may unambiguously extend the function $\iota : (i, a, b)/{\sim_{n'}} \mapsto i$ to both $\mathbf{G}_k$ and $\mathbf{H}_k$. Now define the distance between two elements within $\mathbf{G}_k$ by

$$d_{\mathbf{G}_k}(x, x') := \begin{cases} |\iota(x) - \iota(x')|, & \text{if } x,\, x' \text{ lie in the same block component,} \\ \infty, & \text{otherwise,} \end{cases}$$

with $d_{\mathbf{H}_k}$ defined identically on $\mathbf{H}_k$. The subscripts $\mathbf{G}_k$ or $\mathbf{H}_k$ are dropped when no ambiguity arises. Note that these "distances" are not metrics in the strict sense, as distinct points may have distance 0 (if they lie in the same block component and have the same ι value).

At each round $i \in \{0, 1, \ldots, k\}$, we will say that the distance between two elements x and x' in $\mathbf{G}_k$ or $\mathbf{H}_k$ is large_i if it at least 2^{k-i+1}. Note that this also includes the case where the distance between x and x' is infinite. Assume that Spoiler and Duplicator have played $i' < k$ rounds of the k round game and have selected some elements $g_1, \ldots, g_{i'} \in \mathbf{G}_k$ and $h_1, \ldots, h_{i'} \in \mathbf{H}_k$. Let $x, x' \in \{g_1, \ldots, g_{i'}, h_1, \ldots, h_{i'}\}$. If the distance between x and x' is greater than or equal to $\mathrm{large}_{i'}$ and there are no previously played elements between them, then $2^{k-i'+1}$ is large enough so that Spoiler cannot connect x and x'. More precisely: at each round of play, Spoiler may halve the distance between x and x'; the value of large_i is chosen such that all further plays (to the k^{th} round) will result in a gap of distance at least $\mathrm{large}_k = 2$ somewhere "between" x and x'. See Libkin [11] for very similar arguments.

The "board" on which Spoiler and Duplicator will play the k-round Ehrenfeucht-Fraïssé game essentially consists of lots of copies of $\mathbf{A}$ and $\mathbf{A} \times \mathbf{A}$ (some of which are connected to one another). Of course, the orientation of $\mathbf{A} \times \mathbf{A}$ is interchangeable, because $\mathbf{A} \times \mathbf{A}$ admits a natural automorphism given by $(a, b) \to (b, a)$. However it is convenient for our discussion to fix the ordering on these coordinates, so that we make think of "pinching" the second coordinate on the left side and the first coordinate on the right side. In the same way, $\mathbf{B}_{\mathrm{R}}$ is in fact isomorphic to $\mathbf{B}_{\mathrm{L}}$, but it is convenient for us to visualize them with their pinches in dual orientation.

We can use the following assumption to shed some light on Duplicator's strategy to win the k-round game. We can assume that Spoiler's first move in the k-round Ehrenfeucht-Fraïssé game will be in $\mathbf{P}_{n'}(\mathbf{A})$ in $\mathbf{G}_k$, since $\mathbf{G}_k \backslash \mathbf{P}_{n'}$ is isomorphic to $\mathbf{H}_k$ and so a first move made by Spoiler outside of $\mathbf{P}_{n'}$ is a wasted one. Assume that Spoiler chooses an element g_1 that is distance l strictly less than large_1 to the left pinch in $\mathbf{P}_{n'}$ (the case for the right pinch is symmetric). Since l is small enough so that, after round k, it is possible for Spoiler to "connect" two elements and because the orientations at the left and right pinches differ, Duplicator must choose an element h_1 from $\mathbf{H}_k$ that is a copy of the element g_1 and in the exact corresponding position, that is, exactly distance l from a left pinch. Now assume that Spoiler chooses an element g_1 that is large_1 from both the left and right pinch in $\mathbf{P}_{n'}$. Again, Duplicator is forced to choose an element from a copy of $\mathbf{B}_R$ or $\mathbf{B}_L$ in $\mathbf{H}_k$ and she must select an element h_1 that is also distance large_1 from any right or left pinch in $\mathbf{H}_k$. She can achieve this by selecting an exact corresponding element to g_1 that is somewhere around the middle of a copy of $\mathbf{B}_R$ or $\mathbf{B}_L$ (her

choice of $\mathbf{B}_R$ or $\mathbf{B}_L$ does not matter since they both have diameter greater than or equal to large$_i$).

We will show, inductively, that Duplicator can not only maintain partial isomorphism but also preserve the following conditions at each of the rounds $i \in \{0, \ldots, k\}$. For $0 < l, j < i$,

(1) if $d_{\mathbf{G}_k}(g_l, g_j) < 2^{k-i+1}$, then $d_{\mathbf{H}_k}(h_l, h_j) = d_{\mathbf{G}_k}(g_l, g_j)$;
(2) if $d_{\mathbf{G}_k}(g_l, g_j) \geqslant 2^{k-i+1}$, then $d_{\mathbf{H}_k}(h_l, h_j) \geqslant 2^{k-i+1}$;
(3) if g_j is of distance $l < 2^{k-i+1}$ from a left pinch, then h_j is of distance l from left pinch;
(4) if g_j is of distance $l < 2^{k-i+1}$ from a right pinch, then h_j is of distance l from right pinch;
(5) if g_j is of distance $l < 2^{k-i+1}$ from a left slice, then h_j is of distance l from a left slice;
(6) if g_j is of distance $l < 2^{k-i+1}$ from a right slice, then h_j is of distance l from a right slice;
(7) the dual conditions to the above (in other words, switching the roles of g and h).

The base case holds vacuously. For the induction step, suppose that Duplicator has survived i rounds of the game by maintaining the above conditions to the completion of round i. Assume that Spoiler is making his $(i+1)$st move in $\mathbf{G}_k$ (the case for $\mathbf{H}_k$ is symmetric). If Spoiler's selection for g_{i+1} is equal to some previously played element g_l, where $l < i$, then Duplicator's response should be h_l.

Case 1: Spoiler chooses g_{i+1} so that it has distance greater than or equal to $2^{k-(i+1)+1}$ from any other previously played element in $\mathbf{G}_k$.
Case 1(a): Spoiler chooses g_{i+1} to be distance greater than or equal to $2^{k-(i+1)+1}$ from any pinch or slice in $\mathbf{G}_k$. Duplicator must select h_{i+1} so that it is also distance greater than or equal to $2^{k-(i+1)+1}$ from any other previously played element and from any pinch or slice in $\mathbf{H}_k$. Duplicator can always find a copy of $\mathbf{B}_\mathrm{R}$ or $\mathbf{B}_\mathrm{L}$ in $\mathbf{H}_k$ that contains no previously played points: after round i there are at least $2(k+1) - i$ unplayed copies (each of length 2^{k+1}) remaining in $\mathbf{H}_k$. If Duplicator selects h_{i+1} to be as close as possible to the middle of one of these intervals, then we ensure that the large$_{i+1}$ property and all other conditions of the hypothesis are carried through to the $(i+1)$st round.

Case 1(b): Spoiler chooses g_{i+1} so that it is distance $l < 2^{k-(i+1)+1}$ from a pinch or slice in $\mathbf{G}_k$. Without loss of generality, suppose spoiler chooses g_{i+1} so that it is distance $l < 2^{k-(i+1)+1}$ from a left pinch in $\mathbf{G}_k$. Duplicator must select an element h_{i+1} in $\mathbf{H}_k$ that is exactly distance l from a left pinch in $\mathbf{H}_k$ and such that h_{i+1} is distance greater than or equal to $2^{k-(i+1)+1}$ from any other previously played element. Duplicator can always find a copy of $\mathbf{B}_\mathrm{L}$ in $\mathbf{H}_k$ that contains no previously played elements: after round i there are at least $(k+1) - i$ copies remaining in $\mathbf{H}_k$. Now Duplicator should select an element h_{i+1} in one of these copies corresponding to the element that Spoiler chose in his copy of $\mathbf{B}_\mathrm{L}$ (or $\mathbf{P}_{n'}$) relative to the natural embedding of $\mathbf{B}_\mathrm{L}$ into $\mathbf{P}_{n'}$, and so that it has distance exactly $l < 2^{k-(i+1)+1}$ from the left pinch. Then all the conditions of hypothesis are preserved to the $(i+1)$st round in this case.

Case 2. Spoiler chooses g_{i+1} so that it has distance strictly less than $2^{k-(i+1)+1}$ from some previously played element g_l $(l < i)$ in $\mathbf{G}_k$.

Case 2(a): Spoiler chooses g_{i+1} to be distance greater than or equal to $2^{k-(i+1)+1}$ from any pinch or slice in $\mathbf{G}_k$. Without loss of generality assume that Spoiler has selected g_{i+1} to be to the right of g_l. There are two sub-cases to consider, however Duplicator's strategy is the same in each case: she should select h_{i+1} in $\mathbf{H}_k$ to be an exact copy of the element g_{i+1}, and exactly distance d to the right of h_l. We only need to check that conditions (1) and (2) are maintained in this case.
Case 2(a)(i): If there exists a previously selected element g_j to the right of g_{i+1} that is distance strictly less than $2^{k-(i+1)+1}$, then $d_{\mathbf{G}_k}(g_l, g_j) < 2^{k-i+1}$, and so condition (2) hypothesis tells us that $d_{\mathbf{H}_k}(h_l, h_j) = d_{\mathbf{G}_k}(g_l, g_j)$. Hence Duplicator's selection of h_{i+1} at exactly distance d to the right of h_l ensures that $d_{\mathbf{H}_k}(h_l, h_{i+1}) = d_{\mathbf{G}_k}(g_l, g_{i+1})$ and $d_{\mathbf{H}_k}(h_{i+1}, h_j) = d_{\mathbf{G}_k}(g_{i+1}, g_j)$. Hence all conditions of the hypothesis are maintained in this case.
Case 2(a)(ii): If the distance of any other previously played element to the right of g_{i+1} is greater than or equal $2^{k-(i+1)+1}$. Then, at round i, the distance of any previously played element to the right of g_l was greater than or equal to 2^{k-i+1}, so condition (1) of the hypothesis tells us that the distance of any previously played element to the right of h_l in $\mathbf{H}_k$ is 2^{k-i+1}. Hence Duplicator's selection of h_{i+1} at exactly distance d to the right of h_l ensures that $d_{\mathbf{H}_k}(h_l, h_{i+1}) = d_{\mathbf{G}_k}(g_l, g_{i+1})$ and the distance of any other previously played element to the right of g_{i+1} is greater than or equal $2^{k-(i+1)+1}$. Hence all conditions of the hypothesis are maintained in this case.
Case 2(b) Spoiler chooses g_{i+1} to be distance strictly less than $2^{k-(i+1)+1}$ from any pinch or slice in $\mathbf{G}_k$. Without loss of generality assume that Spoiler chooses g_{i+1} to be distance strictly less than $2^{k-(i+1)+1}$ from a left pinch.
Case 2 (b)(i): If Spoiler chooses g_{i+1} to be to the right of g_l at distance p, then at round i, the distance m between g_l and the left pinch must have been strictly less than 2^{k-i+1}. So condition (3) of the hypothesis ensures that h_l is also distance m from a left pinch. If Duplicator selects h_{i+1} to be distance p to the right of h_l, then h_{i+1} is exactly the same distance $(m+p)$ as g_{i+1} from a left pinch, so condition (3) of the hypothesis is maintained to this round. It remains to check conditions (1) and (2): applying the arguments given in Case 2(a)(i) and Case 2(a)(ii), respectively guarantees that all the conditions of the hypothesis are maintained in this case.
Case 2 (b)(ii): If Spoiler chooses g_{i+1} to be to the left of g_l at distance r strictly less than $2^{k-(i+1)+1}$, then g_l must be distance s strictly less than 2^{k-i+1} from the left pinch, and so the hypothesis ensures that h_l is also distance s from a left pinch. Duplicator should select h_{i+1} to be distance r to the right of h_l, then h_{i+1} is exactly the same distance $(r+s)$ as g_{i+1} from a left pinch, and so condition (3) of the hypothesis is maintained to this round. This selection also guarantees that condition (1) is maintained in this case. If there is no previously played element further to the left of g_{i+1}, then we are done. If on the other hand, there exists a previously played element g_j further to the left of g_{i+1}, we know that $d_{\mathbf{G}_k}(g_j, g_{i+1}) < 2^{k-(i+1)+1}$, and so we need to check condition (1) of the hypothesis. At round i, we had $d_{\mathbf{G}_k}(g_j, g_l) < 2^{k-i+1}$, and so the hypothesis tells us that $d_{\mathbf{H}_k}(h_j, h_l) = d_{\mathbf{G}_k}(g_j, g_l)$, and hence Duplicator's selection ensures that $d_{\mathbf{G}_k}(g_j, g_{i+1}) = d_{\mathbf{H}_k}(h_j, h_{i+1})$. □

Our main result now follows almost immediately.

***Proof of Theorem** 4.8.* We prove the contrapositive. Assume that $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathbf{A})$ is not definable by a finite conjunction of anti-identities modulo $\mathrm{Th}(\mathcal{K})$. Then $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathbf{A})$ fails the finite duality property relative to the class $\mathcal{K}$ (clearly, (1) $\Rightarrow$ (2) of Theorem 4.9 relativises to any class). Hence $\mathrm{CSP}(\mathbf{A})$ fails the finite duality property and so for every natural number n, the structure $\mathbf{P}_n(\mathbf{A})$ does not admit a homomorphism to $\mathbf{A}$. Now since the class $\mathcal{K}$ is closed under forming the n-pinch construction over $\mathbf{A}$ and taking disjoint unions, the proof of (3) $\Rightarrow$ (1) of Theorem 4.9 relativises to $\mathcal{K}$: we have found, for each non-negative integer k, two structures $\mathbf{G}_k$ and $\mathbf{H}_k$ in $\mathcal{K}$ such that:

(1) Duplicator has a winning strategy in the k-round Ehrenfeucht-Fraïssé game played on $\mathbf{G}_k$ and $\mathbf{H}_k$, and
(2) $\mathbf{H}_k$ belongs to $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathcal{A})$, but $\mathbf{G}_k$ does not.

Hence the property of admitting a homomorphism into $\mathbf{A}$ relative to the class $\mathcal{K}$ is not definable by a first-order sentence. □

We now complete this section by extending Theorem 4.8. The proof follows a similar structure: we first prove a non-relativised version, namely Theorem 4.15, from which the extended result will follow almost immediately. The following three results given in Lemma 4.10, Corollary 4.11 and Lemma 4.12 are required. The first lemma appears as Lemma 2.1 in [10], and is proved by a probabilistic argument based on that originally given by Erdős to show the existence, for each non-negative k, of a graph with chromatic number greater than k and no cycles of length less than or equal to k. See the proof given in Theorem 5 of [6].

Lemma 4.10 ([10], Lemma 2.1)**.** *Let $\mathbf{A}$ and $\mathbf{B}$ be $\mathcal{R}$-structures. If there exists no homomorphism from $\mathbf{B}$ to $\mathbf{A}$, then for every natural number n there exists an $\mathcal{R}$-structure $\mathbf{C}_n$ of girth at least n such that the following conditions hold*:

(1) *there exists no homomorphism from $\mathbf{C}_n$ to $\mathbf{A}$*;
(2) *there exists a homomorphism from $\mathbf{C}_n$ to $\mathbf{B}$.*

We observe that Lemma 4.10 is easily extended to a finite number of relational structures.

Corollary 4.11. *Let $m \in \mathbb{N}$, let $\mathcal{F} = \{\mathbf{A}_1, \ldots, \mathbf{A}_m\}$ be a finite set of $\mathcal{R}$-structures and let $\mathbf{B}$ be an $\mathcal{R}$-structure. If $\mathbf{B}$ does not admit a homomorphism to $\mathbf{A}_i$, for each $i \in \{1, \ldots, m\}$, then for any natural number n, there exists an $\mathcal{R}$-structure $\mathbf{C}_n$ of girth at least n such that the following conditions hold*:

(1) *there exists no homomorphism from $\mathbf{C}_n$ to $\mathbf{A}_i$, for each $i \in \{1, \ldots, m\}$*;
(2) *there exists a homomorphism from $\mathbf{C}_n$ to $\mathbf{B}$.*

Proof. Let $i \in \{1, \ldots, m\}$. Since $\mathbf{B}$ does not admit a homomorphism to $\mathbf{A}_i$, we can use Lemma 4.10, to find a structure $\mathbf{C}^i_n$ of girth at least n that admits a homomorphism to $\mathbf{B}$ but does not admit one to $\mathbf{A}_i$. Define $\mathbf{C}_n$ to be the disjoint union of the $\mathbf{C}^i_n$. Clearly, the structure $\mathbf{C}_n$ has girth at least n and does not admit a homomorphism to $\mathbf{A}_i$, for all $i \in \{1, \ldots, m\}$. □

The proof of the next lemma is identical, up to replacing Lemma 2.1 of [10] with its extended version Corollary 4.11, to the proof of Lemma 2.4 in [10] . We refer the reader to the proof there.

Lemma 4.12. *Let $\mathcal{F}$ be a finite set of homomorphism-independent $\mathcal{R}$-structures. Then the following are equivalent*:

(1) $\mathbb{H}^{\leftarrow}_{\text{fin}}(\mathfrak{F})$ *satisfies the finite duality property*;
(2) *there is an upper bound on the diameter of the critical obstructions for* $\mathbb{H}^{\leftarrow}_{\text{fin}}(\mathfrak{F})$.

Part of the proof of Proposition 4.1 in [10] is applicable to a slightly more general setting. We state this as a lemma after first fixing some useful notation.

Notation 4.13. Let $\mathfrak{F}$ be a class of finite $\mathcal{R}$-structures, let $\mathbf{C}$ be a critical obstruction for the class $\mathbb{H}^{\leftarrow}_{\text{fin}}(\mathfrak{F})$ and let $x \in C$. We let $\mathbf{C}_x$ denote the substructure of $\mathbf{C}$ induced by the set $C\backslash\{x\}$.

Lemma 4.14. *Let* $n, m \in \mathbb{N}$*, let* $\mathfrak{F} = \{\mathbf{A}_1, \ldots, \mathbf{A}_m\}$ *be a finite set of homomorphism-independent* $\mathcal{R}$*-structures and let* $\mathbf{C}$ *be a critical obstruction for the class* $\mathbb{H}^{\leftarrow}_{\text{fin}}(\mathfrak{F})$ *with diameter at least* $n+2$*. Assume that there are* $x, y \in C$ *and* $j \in \{1, \ldots, m\}$ *with* x *and* y *distance* $n+2$ *apart and with both* $\mathbf{C}_x$ *and* $\mathbf{C}_y$ *admitting a homomorphism to* $\mathbf{A}_j$*. Then there exists a homomorphism from* $\mathbf{C}$ *to* $\mathbf{P}_n(\mathbf{A}_j)$*.*

Proof. The lemma is essentially borrowed from [10], so we give only a proof sketch. Let x and y be elements in $\mathbf{C}$ at distance $n+2$ and let $\alpha : \mathbf{C}_x \to \mathbf{A}_j$ and $\beta : \mathbf{C}_y \to \mathbf{A}_j$ be homomorphisms. The maps κ and ϕ defined below are shown to be homomorphisms in the proof of Proposition 4.1 in [10]. Define $\kappa : C \to L_n$ by

$$\kappa(z) := \begin{cases} 0, & \text{if } z = x, \\ d_{\mathbf{C}}(x,z) - 1, & \text{if } d_{\mathbf{C}}(x,z) \leq n+1 \text{ and } z \neq x, \\ n, & \text{if } d_{\mathbf{C}}(x,z) \geq n+2 \end{cases}$$

Now select an element $a \in A$ and define the map ϕ from C to $P_n(\mathbf{A}_j)$ by

$$\phi(z) := \begin{cases} ((\kappa(z), \alpha(z), \beta(z))/\sim_n, & \text{if } z \notin \{x, y\}, \\ ((\kappa(z), \alpha(z), a)/\sim_n, & \text{if } z = x, \\ ((\kappa(z), a, \beta(z))/\sim_n, & \text{if } z = y. \end{cases}$$

This is the desired homomorphism $\phi : \mathbf{C} \to \mathbf{P}_n(\mathbf{A}_j)$. □

We need one more result before we can prove our main result for this section. Recall first that $\mathbb{H}^{\leftarrow}_{\text{fin}}(\mathfrak{F}) = \bigcup_{i \in \{1,\ldots,m\}} \text{CSP}(\mathbf{A}_i)$.

Theorem 4.15. *Let* $m \in \mathbb{N}$ *and let* $\mathfrak{F} = \{\mathbf{A}_1, \ldots, \mathbf{A}_m\}$ *be a finite set of homomorphism-independent* $\mathcal{R}$*-structures. Then the following are equivalent*:

(1) $\mathbb{H}^{\leftarrow}_{\text{fin}}(\mathfrak{F})$ *is definable by a first-order sentence*;
(2) *For each* $i \in \{1, \ldots, m\}$*, the class* $\text{CSP}(\mathbf{A}_i)$ *satisfies the finite duality property*;
(3) $\mathbb{H}^{\leftarrow}_{\text{fin}}(\mathfrak{F})$ *satisfies the finite duality property*;
(4) $\mathbb{H}^{\leftarrow}_{\text{fin}}(\mathfrak{F})$ *is definable by a finite conjunction of anti-identities.*

Proof. (1) ⇒ (2): We prove the contrapositive. Assume that $\text{CSP}(\mathbf{A}_j)$ fails to satisfy the finite duality property. For each $k \in \mathbb{N} \cup \{0\}$, define $n_k := 2^{k+1} + 1$ and let $\mathbf{P}_{n_k}(\mathbf{A}_j)$ be the $(2^{k+1}+1)$-pinch over $\mathbf{A}_j$. Let $\mathbf{B}_R$ and $\mathbf{B}_L$ be the substructures of $\mathbf{P}_{n_k}(\mathbf{A}_j)$ induced by the sets $B_R = \{(k, a, b)/\sim_{n'} |\ k \neq 0\}$ and $B_L = \{(k, a, b)/\sim_{n'} |\ k \neq n'\}$, respectively. From part (1 and 2) of Theorem 4.5, we have that, for all $k \in \mathbb{N}$, the structure $\mathbf{P}_{n_k}(\mathbf{A}_j)$ does not admit a homomorphism to $\mathbf{A}_j$, but the substructures $\mathbf{B}_R$ and $\mathbf{B}_L$ both admit homomorphisms to $\mathbf{A}_j$. Now consider the disjoint union $\mathbf{P}_{n_k}(\mathbf{A}_j) \dot{\cup} \mathbf{A}_j$. Since $\mathfrak{F} = \{\mathbf{A}_1, \ldots, \mathbf{A}_m\}$ is a

homomorphism independent set, it follows that, for all $k \in \mathbb{N} \cup \{0\}$, the structure $\mathbf{P}_{n_k}(\mathbf{A}_j) \dot\cup \mathbf{A}_j$ does not belong to the class $\mathbb{H}^{\leftarrow}_{\mathrm{fin}}(\mathcal{F})$. For each $k \in \mathbb{N} \cup \{0\}$ let $\mathbf{H}_k$ be the structure obtained by taking the disjoint union of $\mathbf{A}_j$ with $k+1$ copies of $\mathbf{B}_R \dot\cup \mathbf{B}_L$, and let $\mathbf{G}_k = \mathbf{P}_{n'}(\mathbf{A}_j) \dot\cup \mathbf{H}_k$. Then, for each $k \in \mathbb{N} \cup \{0\}$, we have $\mathbf{H}_k \in \mathrm{CSP}(\mathbf{A}_j) \subseteq \mathbb{H}^{\leftarrow}_{\mathrm{fin}}(\mathcal{F})$, while $\mathbf{G}_k \notin \mathbb{H}^{\leftarrow}_{\mathrm{fin}}(\mathcal{F})$. Now apply the argument given in Theorem 4.8 with the extra condition that whenever Spoiler selects an element g from $\mathbf{G}_k$ in the extra copy of $\mathbf{A}_j$, then Duplicator selects the corresponding element h from the extra copy of $\mathbf{A}_j$ in $\mathbf{H}_k$, and vice-versa. Hence, for each non-negative integer k, Duplicator has a winning strategy in the k-round Ehrenfeucht-Fraïssé game played on $\mathbf{G}_k$ and $\mathbf{H}_k$, and so the property of admitting a homomorphism to $\mathcal{F}$ is not definable by a first-order sentence.

(2) ⇒ (3): We will show that there is a bound on the the diameter of the critical obstructions for $\mathbb{H}^{\leftarrow}_{\mathrm{fin}}(\mathcal{F})$. The required result will then follow by Lemma 4.12. Assume that, for each $i \in \{1, \ldots, m\}$, the class $\mathrm{CSP}(\mathbf{A}_i)$ satisfies the finite duality property. Then, for each $i \in \{1, \ldots, n\}$, there exists a natural number l_i with $\mathbf{P}_{l_i}(\mathbf{A}_i)$ admitting a homomorphism to $\mathbf{A}_i$, by part (2) of Theorem 4.5. Let $\mathbf{C}$ be any critical obstruction for $\mathbb{H}^{\leftarrow}_{\mathrm{fin}}(\mathcal{F})$. Of course, for any element $z \in C$, the substructure $\mathbf{C}_z$ admits a homomorphism to at least one of the $\mathbf{A}_i$. Now let k be the diameter of $\mathbf{C}$, and let x and y be elements in C at distance k. We need to consider two cases.

Case 1: there exists $j \in \{1, \ldots, m\}$ with $\mathbf{C}_x$ and $\mathbf{C}_y$ admitting homomorphisms to $\mathbf{A}_j$. Lemma 4.14 tells us that $\mathbf{C}$ admits a homomorphism to $\mathbf{P}_{k-2}(\mathbf{A}_j)$. But since $\mathbf{P}_{l_j}(\mathbf{A}_j)$ admits a homomorphism to $\mathbf{A}_j$, and therefore $\mathbf{P}_l(\mathbf{A}_j)$ admits a homomorphism to $\mathbf{A}_j$, for all $l \geq l_j$ (by part 3 of Theorem 4.5), we must have that k is strictly less than $l_j + 2$.

Case 2: there exists $p, q \in \{1, \ldots, m\}$ with $p \neq q$ and with $\mathbf{C}_x$ admitting a homomorphism to $\mathbf{A}_p$ and $\mathbf{C}_y$ admittng a homomorphism to $\mathbf{A}_q$. We will show that k is less than or equal to $l_1 + l_2 + \cdots + l_m + 2m$. If there exists $x', y' \in C$ and $j \in \{1, \ldots, m\}$ with $\mathbf{C}_{x'}$ and $\mathbf{C}_{y'}$ admitting homomorphisms to A_j, then from Case 1, the maximum distance between x' and y' in $\mathbf{C}$ is strictly less than $l_j + 2$. Thus, since $\mathbf{C}_x$ admits a homomorphism to $\mathbf{A}_p$, any element $z \in C$ with $\mathbf{C}_z$ admitting a homomorphism to $\mathbf{A}_p$ must be at most distance $l_p + 1$ from x. Therefore any element $w \in C$ with distance exactly $l_p + 2$ from x must have $\mathbf{C}_w$ admitting a homomorphism to $\mathbf{A}_s$, where $s \in \{1, \ldots, m\}$ and $s \neq p$. Then any element $c \in C$ with distance $l_s + 2$ from w (and distance $l_p + l_s + 4$ from x) must have $\mathbf{C}_c$ admitting a homomorphism to $\mathbf{A}_t$, for some $t \notin \{p, s\}$. We can continue this process until we run out of $\mathbf{A}_i$'s. It follows that the diameter k of $\mathbf{C}$ is at most $l_1 + l_2 + \cdots + l_m + 2m$.

The implication (3) ⇒ (4) is routine, see Lemma 2 in [9] for example, and the implication (4) ⇒ (1) is trivial. □

Now the proof of our extended result follows almost immediately.

Theorem 4.16. *Let $m \in \mathbb{N}$. Let $\mathcal{K}$ be a class of finite relational structures in the signature $\mathcal{R}$ and let $\mathcal{F} = \{\mathbf{A}_1, \ldots, \mathbf{A}_m\}$ be a finite set of homomorphism-independent structures in $\mathcal{K}$. If $\mathcal{K}$ is closed under forming the n-pinch construction over $\mathcal{F}$ and taking disjoint unions, then $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathcal{F})$ is equal to $\mathcal{K} \cap \mathrm{Mod}(\phi)$ for some first order sentence ϕ if and only if there is a finite conjunction of anti-identities ψ such that $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathcal{F}) = \mathcal{K} \cap \mathrm{Mod}(\psi)$.*

Proof. We prove the contrapositive. Assume that $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathfrak{F})$ is not definable by a finite conjunction of anti-identities modulo $\mathrm{Th}(\mathcal{K})$. Then $\mathbb{H}^{\leftarrow}_{\mathcal{K}}(\mathfrak{F})$ fails the finite duality property relative to the class $\mathcal{K}$, and therefore $\mathbb{H}_{\mathrm{fin}}(\mathfrak{F})$ fails the finite duality property. So there is some $j \in \{1, \dots, m\}$ with $\mathrm{CSP}(\mathbf{A}_j)$ failing to satisfy the finite duality property, using (2) $\Rightarrow$ (3) of Theorem 4.15. Now since the class $\mathcal{K}$ is closed under forming the n-pinch construction over $\mathfrak{F}$ and taking disjoint unions, the proof of (1) $\Rightarrow$ (2) of Theorem 4.15 relativises to $\mathcal{K}$. Hence the property of admitting a homomorphism into $\mathfrak{F}$ relative to the class $\mathcal{K}$ is not definable by a first-order sentence. □

We have made multiple references to closure under formation of the n-pinch and taking disjoint unions. We conclude this section with some examples of sentences that are closed under these constructions.

Example 4.17. *Let $\mathcal{R} = \{R_1, \dots, R_m\}$ be a relational signature. Then a class of $\mathcal{R}$-structures $\mathfrak{F}$ is closed under forming the n-pinch (for sufficiently large n) and disjoint unions if it is definable by sentences of any of the following forms*:

(1) reflexivity: $\forall x\, (x, \dots, x) \in R_i$;
(2) any finite system of anti-identities;
(3) implications of the form

$$(*) \qquad \forall x_1, \dots, x_{r_i}\ [(x_1, \dots, x_{r_i}) \in R_i \to (x_{l_1}, \dots, x_{l_{r_j}}) \in R_j],$$

where $\{l_1, \dots, l_{r_j}\} \subseteq \{1, \dots, r_i\}$ and $|\{x_1, \dots, x_{r_i}\}| = r_i$.

Proof. Closure under disjoint unions is obvious in each case; in case (3) this uses the fact that the conclusion of the implication involves variables only appearing in the premise. We now discuss the formation of n-pinches (for sufficiently large n).

Case (1) is trivial, while case (2) follows from the fact that for $m < n$, the m-vertex substructures of the n-pinch over $\mathbf{A}$ admit a homomorphism into $\mathbf{A}$ (by Theorem 4.5(1)), and hence the n-pinch will satisfy any anti-identity in at most m variables that is true on $\mathbf{A}$.

For Case (3), let $(x_1, \dots, x_{r_i}) = ((k_1, a_1, b_1)/{\sim_n}, \dots, (k_{r_i}, a_{r_i}, b_{r_i})/{\sim_n}) \in R_i^{\mathbf{P}_n(\mathbf{A})}$. We have $\{k_1, \dots, k_{r_i}\} \subseteq \{k, k+1\}$. If $k = 0$, then from the definition of $\sim_n$, for each $m \in \{1, \dots, r_i\}$, there exists $(k_m, a_m, b'_m) \in (k_m, a_m, b_m)/{\sim_n}$ satisfying $((k_1, a_1, b'_1), \dots, (k_{r_i}, a_{r_i}, b'_{r_i})) \in R_i^{\mathbf{L}_n \times \mathbf{A} \times \mathbf{A}}$. It follows that $(a_1, \dots, a_{r_i}) \in R_i^{\mathbf{A}}$ and $(b'_1, \dots, b'_{r_i}) \in R_i^{\mathbf{A}}$. Note that here we are implicitly using the condition that $|\{x_1, \dots, x_{r_i}\}| = r_i$; if $x_1, \dots, x_{r_i}$ contained repeated elements then we could not guarantee repeated elements in $b'_1, \dots, b'_{r_i}$. Now interpreting $(*)$ in $\mathbf{A}$, gives $(a_{l_1}, \dots, a_{l_{r_j}}) \in R_j^{\mathbf{A}}$ and $(b'_{l_1}, \dots, b'_{l_{r_j}}) \in R_j^{\mathbf{A}}$, where $\{l_1, \dots, l_{r_j}\} \subseteq \{1, \dots, r_i\}$. It follows that $(x_{l_1}, \dots, x_{l_{r_j}}) \in R_j^{\mathbf{P}_n(\mathbf{A})}$. A similar argument works for $k = n$ and even easier argument works for $k \in \{1, \dots, n-1\}$. □

The class of simple graphs for example is definable by one anti-identity $(\forall x) \neg R(x, x)$ and one implication $(\forall x, y)\ R(x, y) \to R(y, x)$ of the form described in Example 4.17(3).

References

[1] M. Ajtai and Y. Gurevich. Monotone versus positive. *J. Assoc. Comput. Mach.*, 34(4):1004–1015, 1987.

[2] A. Atserias. On digraph coloring problems and treewidth duality. *European J. Combin.*, 29(4):796–820, 2008.
[3] A. Atserias, A. Dawar, and M. Grohe. Preservation under extensions on well-behaved finite structures. *SIAM J. Comput.*, 38(4):1364–1381, 2008.
[4] A. Atserias, A. Dawar, and P. G. Kolaitis. On preservation under homomorphisms and unions of conjunctive queries. *J. ACM*, 53(2):208–237 (electronic), 2006.
[5] D. M. Clark, B. A. Davey, M. G. Jackson, and J. G. Pitkethly. The axiomatizability of topological prevarieties. *Adv. Math.*, 218(5):1604–1653, 2008.
[6] T. Feder and M. Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory. *SIAM J. Comput.*, 28(1):57–104 (electronic), 1999.
[7] V. A. Gorbunov. *Algebraic Theory of Quasivarieties*. Siberian School of Algebra and Logic. Consultants Bureau, New York, 1998. Translated from the Russian.
[8] Y. Gurevich. Toward logic tailored for computational complexity. In *Computation and proof theory (Aachen, 1983)*, volume 1104 of *Lecture Notes in Math.*, pages 175–216. Springer, Berlin, 1984.
[9] M. Jackson and B. Trotta. Constraint satisfaction, irredundant axiomatisability and continuous colouring. *Studia Logica*, 101(1):65–94, 2013.
[10] B. Larose, C. Loten, and C. Tardif. A characterisation of first-order constraint satisfaction problems. *Log. Methods Comput. Sci.*, 3(4):4:6, 22, 2007.
[11] L. Libkin. *Elements of Finite Model Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer-Verlag, Berlin, 2004.
[12] B. Rossman. Homomorphism preservation theorems. *J. ACM*, 55(3):Art. 15, 53, 2008.
[13] W. W. Tait. A counterexample to a conjecture of Scott and Suppes. *J. Symb. Logic*, 24:15–16, 1959.

Department of Mathematics and Statistics, La Trobe University, Victoria 3086, Australia

E-mail address: leham@students.latrobe.edu.au